

A minimum-change version of the Chung-Feller theorem for Dyck paths

Torsten Mütze, Christoph Standke, Veit Wiechert

Institut für Mathematik

TU Berlin, 10623 Berlin, Germany

{muetze, standke, wiechert}@math.tu-berlin.de

ABSTRACT. A Dyck path with $2k$ steps and e flaws is a path in the integer lattice that starts at the origin and consists of k many \nearrow -steps and k many \searrow -steps that change the current coordinate by $(1, 1)$ or $(1, -1)$, respectively, and that has exactly e many \searrow -steps below the line $y = 0$. Denoting by D_{2k}^e the set of Dyck paths with $2k$ steps and e flaws, the Chung-Feller theorem asserts that the sets $D_{2k}^0, D_{2k}^1, \dots, D_{2k}^k$ all have the same cardinality $\frac{1}{k+1} \binom{2k}{k} = C_k$, the k -th Catalan number. The standard combinatorial proof of this classical result establishes a bijection f' between D_{2k}^e and D_{2k}^{e+1} that swaps certain parts of the given Dyck path x , with the effect that x and $f'(x)$ may differ in many positions. In this paper we strengthen the Chung-Feller theorem by presenting a simple bijection f between D_{2k}^e and D_{2k}^{e+1} which has the additional feature that x and $f(x)$ differ in only *two* positions (the least possible number). We also present an algorithm that allows to compute a sequence of applications of f in constant time per generated Dyck path. As an application, we use our minimum-change bijection f to construct cycle-factors in the odd graph O_{2k+1} and the middle levels graph M_{2k+1} — two intensively studied families of vertex-transitive graphs — that consist of C_k many cycles of the same length.

1. INTRODUCTION

The Catalan numbers are one of the most fundamental counting sequences in combinatorics, and Dyck paths are among the most heavily studied Catalan families (see [Sta15] and references therein). A *Dyck path with $2k$ steps and e flaws* ($e \in \{0, 1, \dots, k\}$) is a path in the integer lattice \mathbb{Z}^2 that starts at the origin $(0, 0)$ and consists of k many \nearrow -steps and k many \searrow -steps that change the current coordinate by $(1, 1)$ or $(1, -1)$, respectively (so the path ends at the coordinate $(2k, 0)$), and that has exactly e many \searrow -steps below the line $y = 0$. In particular, a Dyck path with zero flaws never moves below the line $y = 0$, and a Dyck path with k flaws never moves above the line $y = 0$. We denote the set of all Dyck paths with $2k$ steps and e flaws by D_{2k}^e . Clearly, the total number of Dyck paths with $2k$ steps and an arbitrary number of flaws is $\binom{2k}{k}$. The classical Chung-Feller theorem, first proved in [Mac09, p. 168], asserts that partitioning the set of all these Dyck paths according to the number of flaws yields $k + 1$ subsets of the same size $\frac{1}{k+1} \binom{2k}{k} =: C_k$, where C_k is the k -th Catalan number.

Theorem 1 ([Mac09], [CF49]). *For any $k \geq 1$ we have $|D_{2k}^0| = |D_{2k}^1| = \dots = |D_{2k}^k| = \frac{1}{k+1} \binom{2k}{k} = C_k$.*

Since its discovery, several proofs and generalizations of Theorem 1 have appeared in the literature [Mac09, CF49, Hod55, Nar67, Che08, Ruk11, Woa01, EFY05, MY09]. The standard combinatorial proof, first presented in [Hod55] and rediscovered in [Che08], establishes the following bijection f' between the sets D_{2k}^e and D_{2k}^{e+1} (see Figure 1): For any Dyck path $x \in D_{2k}^e$, considering the first \nearrow -step above the line $y = 0$ and the first \searrow -step returning to the line $y = 0$, and denoting the corresponding indices by a and b , the path $x = (x_1, x_2, \dots, x_{2k})$ can be partitioned as $x =$




FIGURE 1. Bijection f' used to prove Theorem 1. The Dyck path x has 4 flaws, the Dyck path $f'(x)$ has 5 flaws.




FIGURE 2. Definition of the bijection f from Theorem 2. The Dyck path x has 4 flaws, the Dyck path $f(x) = h(g(x))$ has 5 flaws, and x and $f(x)$ differ in only two positions marked with gray vertical bars. In the upper part of the figure, all \searrow -steps of x touching the line $y = 0$ are drawn gray. In the middle part, all \nearrow -steps of x' touching the line $y = 1$ are drawn gray.

$u \circ \nearrow \circ v \circ \searrow \circ w$ where $u := (x_1, \dots, x_{a-1})$, $v := (x_{a+1}, \dots, x_{b-1})$ and $w := (x_{b+1}, \dots, x_{2k})$, and \circ denotes the concatenation. Then the image of x under f' is defined as $f'(x) := v \circ \searrow \circ u \circ \nearrow \circ w$. It is easy to check that f' increases the number of flaws by 1 and that it is indeed a bijection between D_{2k}^e and D_{2k}^{e+1} (see Figure 3). Note however that the number of positions in which $x = u \circ \nearrow \circ v \circ \searrow \circ w$ and $f'(x) = v \circ \searrow \circ u \circ \nearrow \circ w$ differ can be large due to the swapping of the subpaths u and v . In the worst case the two Dyck paths may differ in every position, see e.g. the paths x' and $f'(x')$ on the right-hand side of Figure 3.

1.1. Our results. The main result of this paper is the following strengthening of Theorem 1: We establish a bijection f between D_{2k}^e and D_{2k}^{e+1} such that for any $x \in D_{2k}^e$, the Dyck paths x and $f(x)$ differ in only *two* positions (two changes are clearly the least possible number). In other words: $f(x)$ is obtained from x by flipping a single \nearrow -step and a single \searrow -step, leaving the initial and terminal parts of the path unchanged, and shifting down the middle part by two units.




FIGURE 3. Illustration of the bijection f' used to prove Theorem 1 on the set of all Dyck paths of length 6.




FIGURE 4. Illustration of the minimum-change bijection f from Theorem 2 on the set of all Dyck paths of length 6. The gray bars highlight the two positions that are flipped in each step, and the sequences at the bottom the entire sequence of positions that are flipped when repeatedly applying f . The numbers 1–20 and i–v, respectively, indicate related minimum-change orderings of all Dyck paths (with an arbitrary number of flaws) and those with zero flaws as explained in Section 1.1.1.

This minimum-change bijection f is defined as follows (see Figure 2): For any lattice path x and any $c \in \mathbb{Z}$, we let $u_c(x)$ and $d_c(x)$ denote the number of \nearrow -steps or \searrow -steps, respectively, of x starting on the line $y = c$. We say that an \nearrow -step or \searrow -step of x *touches* the line $y = c$, if it starts or ends on this line. For any $x \in D_{2k}^e$, $e < k$, we let $x' := g(x)$ denote the lattice path obtained from x by flipping the $(d_0(x) + 1)$ -th \searrow -step of x touching the line $y = 0$. Similarly, we let $h(x')$ denote the lattice path obtained from x' by flipping the $u_1(x')$ -th \nearrow -step of x' touching the line $y = 1$. We then define $f(x) := h(x') = h(g(x))$. Figure 4 shows how f acts on the set of all Dyck paths of length 6 (the two flipped positions are highlighted with gray vertical bars).

Theorem 2. *For any $k \geq 1$ and any $e \in \{0, 1, \dots, k-1\}$ the mapping f is a bijection between D_{2k}^e and D_{2k}^{e+1} such that for any $x \in D_{2k}^e$, the Dyck paths x and $f(x)$ differ in exactly two positions.*

1.1.1. *Efficient algorithms.* We also present an algorithm that allows to compute repeated applications of the minimum-change bijection f from Theorem 2 efficiently.

Theorem 3. *There is an algorithm that computes for any given $k \geq 1$, $e \in \{0, 1, \dots, k-1\}$ and $x \in D_{2k}^e$ each of the Dyck paths $f(x), f^2(x), \dots, f^{k-e}(x)$ in time $\mathcal{O}(1)$ and space $\mathcal{O}(k)$ (overall time including initialization is $\mathcal{O}(k)$).*

Being able to generate each Dyck path in constant time is clearly best possible. We implemented this algorithm in C++, and we invite the reader to experiment with this code, which can be downloaded from the authors' websites [www].

Under this algorithmic lense, the minimum-change bijection f can be seen as a *combinatorial Gray code*. This term refers to the problem of efficiently generating all objects in a particular combinatorial class — e.g. permutations, partitions, subsets, strings, trees, Dyck paths etc. — which is a fundamental task in the area of combinatorial algorithms, as such generation algorithms are used as building blocks in a wide range of practical applications (the survey [Sav97] lists numerous references, see also [NW75, Wil89]). Let us add a few remarks and pointers to related work on combinatorial Gray codes that put the results of this paper into perspective. Gray codes are named after Frank Gray, who invented a method to generate all 2^n subsets of $\{1, 2, \dots, n\}$ such that any two consecutive subsets differ by adding or removing a single element [Gra] (see also [Knu11, Section 7.2.1.1]). Clearly, this problem is equivalent to generating all 2^n bitstrings of length n such that any two consecutive bitstrings differ in exactly one bit. A closely related problem that has received considerable attention in the literature is to generate all subsets of $\{1, 2, \dots, n\}$ of size exactly k for some $k \leq n$, such that any two consecutive subsets differ by exchanging a single element [TL73, Ehr73, BER76] (see also [Knu11, Section 7.2.1.3]). Observe that this is equivalent to generating all bitstrings of length n with exactly k entries equal to 1 such that any two consecutive bitstrings differ in exactly two bits (a 0-bit and a 1-bit are swapped). A special case are the Dyck paths of length $2k$ (with an arbitrary number of flaws) considered in this paper, which can be interpreted as a subset of $\{1, 2, \dots, 2k\}$ of size k or a bitstring of length $2k$ with k entries equal to 1 (every \nearrow -step corresponds to a 1-bit, and every \searrow -step to a 0-bit). Combinatorial Gray codes for these (n, k) -combinations can be computed efficiently even if we impose further constraints on the allowed swaps, e.g. we may only allow swaps of a 0-bit and 1-bit when all bits in between are equal to 0 [EM84]. Even more restrictively, we may only allow swaps of bits that are at most 2 positions apart [Cha89, JM95], or only between adjacent bits [EHR84, Rus88]. In the latter case n and k are required to be even and odd, respectively, or $k \in \{0, 1, n-1, n\}$. In particular, for odd k we may generate all Dyck paths of length $2k$ (with an arbitrary number of flaws) using only such adjacent swaps (a valley $\nearrow\searrow$ becomes a peak $\searrow\nearrow$ or vice versa; such an ordering 1–20 is indicated in Figure 4 for $k = 3$).

We now restrict our attention to Dyck paths with $2k$ steps and zero flaws, which via standard Catalan bijections can also be viewed as well-formed parentheses expressions consisting of k pairs of parentheses, (rooted) ordered trees with k edges, (rooted) binary trees with k vertices, triangulations of a convex $(k+2)$ -gon etc. (see Figure 5, all the bijections between these classes and many more are treated in depth in [Sta15]). Constant average time algorithms for generating all Dyck paths in D_{2k}^0 such that any two consecutive paths differ in exactly two positions have been developed in [RP90, Wal98]. As before, we may even impose the restriction of only allowing swaps of a 0-bit and 1-bit when all bits in between are equal to 0, or only swaps between adjacent bits [BR98], where in the latter case k is required to be even or strictly less than 5 (such an ordering with roman numbers i–v is indicated in Figure 4 for $k = 3$). Summarizing, the minimum-change bijection f




FIGURE 5. Representation of (a) Dyck paths with $2k$ steps and zero flaws for $k = 3$ as (b) well-formed parentheses expressions consisting of k pairs of parentheses, (c) ordered trees with k edges, (d) binary trees with k vertices, (e) triangulations of a convex $(k + 2)$ -gon.

from Theorems 2 and 3 nicely complements the minimum-change generation algorithms for Dyck paths (or parentheses expressions, ordered trees, binary trees, triangulations etc.) known from the literature. In a sense, our bijection is orthogonal to the combinatorial Gray codes mentioned before and indicated in Figure 4 (the bijection f operates on the columns in this figure). We will show that for each $x \in D_{2k}^0$, the Dyck path $f^k(x)$ is obtained by mirroring x at the horizontal line $y = 0$ (see Figure 4, where the Dyck paths at the top and bottom of each column are mirror images of each other). It follows that our bijection can be combined with the before-mentioned algorithms that generate the Dyck paths in D_{2k}^0 and those in D_{2k}^k to obtain new combinatorial Gray codes that enumerate all Dyck paths (with an arbitrary number of flaws) by flipping two positions at a time. Such a Gray code is obtained by iterating through columns and rows in Figure 4 in a saw-tooth fashion. Specifically, f and f^{-1} are applied to move along the columns, and the before-mentioned generation algorithms for D_{2k}^0 and D_{2k}^k are used to move along the first and last row. Similarly, we obtain new combinatorial Gray codes that enumerate all the Dyck paths in D_{2k}^e for some $e \in \{0, 1, \dots, k\}$ by flipping $4 \cdot \min\{e, k - e\} + 2$ positions at a time (again, we move through Figure 4 in a saw-tooth fashion).

1.1.2. Applications. We conclude this section by presenting two concrete applications of Theorem 2. To discuss them we consider the *odd graph* O_{2k+1} , defined for any integer $k \geq 1$ as the graph that has as vertices all k -element subsets of $\{1, 2, \dots, 2k + 1\}$, with an edge between any two sets that are disjoint. Similarly, the *middle levels graph* M_{2k+1} has as vertices all k -element and all $(k + 1)$ -element subsets of $\{1, 2, \dots, 2k + 1\}$, with an edge between any two sets where one is a subset of the other (see Figure 6).

The odd graph and the middle levels graph are families of vertex-transitive graphs that have long been conjectured to have a Hamilton cycle for every value of $k \geq 1$ with one notable exception, the Petersen graph O_5 (these conjectures originated in [ML72, Big79] and [Hav83, BW84], respectively, and they are mentioned e.g. in the popular books [Win04, Knu11, DG12]). Despite the simple definition and despite considerable efforts of many researchers during that last 40 years (see [Mat76, HW78, CL87, Che00, CF02, Che03, Joh04, Joh11] and [Sav93, FT95, SW95, Joh04, HKRR05, GŠ10, SSS09, SA11], respectively), these conjectures remained open for very long (a detailed account of the historic developments can be found in [MS15]). Only recently, it has been shown that the middle

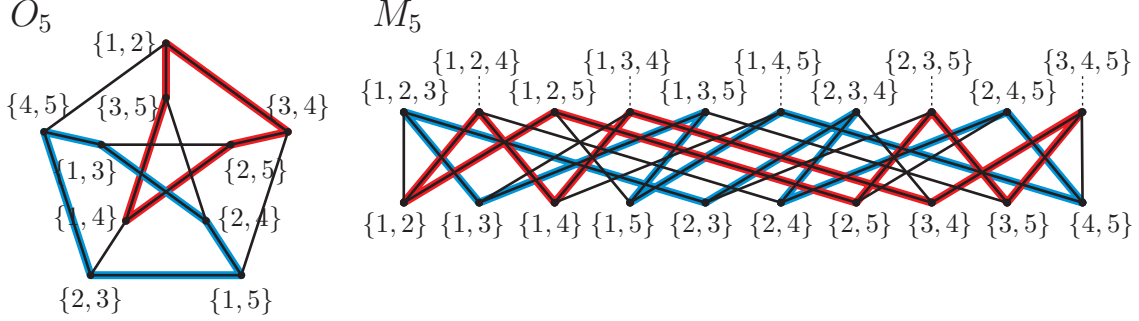


FIGURE 6. Illustration of the odd graph O_{2k+1} (left) and the middle levels graph M_{2k+1} (right) for $k = 2$, and a \mathcal{C}_{2k+1} -factor and a \mathcal{C}_{4k+2} -factor in these graphs (each consisting of two cycles, one colored red and one colored blue), respectively, constructed from our minimum-change bijection f .

levels graph M_{2k+1} indeed has a Hamilton cycle for every $k \geq 1$ [Müt16] (see also [MN15, MN16]). Complementing earlier results [KT88, DSW88, DKS94, JK04, MW12, CGH13] in this direction, we ask more generally, for which divisors ℓ of $n := |V(G)|$ does one of these graphs G have a \mathcal{C}_ℓ -factor, i.e., a collection of n/ℓ many vertex-disjoint cycles, each of length ℓ ($V(G)$ denotes the vertex set of G)? It turns out that the only general case where a positive answer to this question is known is the before-mentioned Hamiltonicity of the middle levels graph (the case $\ell = n$). We use the minimum-change bijection f from Theorem 2 to produce two more positive answers (for values $\ell < n$).

Theorem 4. *For any $k \geq 1$, the odd graph O_{2k+1} has a \mathcal{C}_{2k+1} -factor.*

Theorem 5. *For any $k \geq 1$, the middle levels graph M_{2k+1} has a \mathcal{C}_{4k+2} -factor.*

Note that the \mathcal{C}_{2k+1} -factor of the odd graph and the \mathcal{C}_{4k+2} -factor of the middle levels graph each consist of C_k (the k -th Catalan number) many cycles (see Figure 6). It follows that these cycle-factors have exponentially (in k) many cycles of linear length. It remains an open problem whether these cycles can be joined in a straightforward way to e.g. a Hamilton cycle.

2. PROOFS

In this section we present the proofs of Theorems 2–5. We begin by establishing several lemmas that capture various key properties of the mappings g and h defined in Section 1.1 (recall that the minimum-change bijection f is the composition of g and h).

2.1. Combinatorial properties of the mappings g, h and f . In the following we consider the set $L_{2k,k}$ of all lattice paths with $2k$ steps and exactly k many \nearrow -steps. These paths end at the point $(2k, 0)$, and we have $L_{2k,k} = D_{2k}^0 \cup D_{2k}^1 \cup \dots \cup D_{2k}^k$. We also consider the set $L_{2k,k+1}$ of all lattice paths with $2k$ steps and exactly $k+1$ many \nearrow -steps. These paths clearly end at the point $(2k, 2)$. Our proofs deal with four bijections between the following sets of lattice paths

$$\begin{aligned} g &: L_{2k,k} \setminus D_{2k}^k \rightarrow L_{2k,k+1} \quad , \\ g' &: L_{2k,k} \setminus D_{2k}^0 \rightarrow L_{2k,k+1} \quad , \\ h &: L_{2k,k+1} \rightarrow L_{2k,k} \setminus D_{2k}^0 \quad , \\ h' &: L_{2k,k+1} \rightarrow L_{2k,k} \setminus D_{2k}^k \quad . \end{aligned}$$

Each of these bijections flips exactly one step of the lattice path that is provided as a function argument: g and g' replace a \searrow -step by an \nearrow -step, and h and h' replace an \nearrow -step by a \searrow -step. Recall the definitions of the mappings g and h from Section 1.1. The mappings g' and h' are defined very similarly as follows: For any $x \in L_{2k,k} \setminus D_{2k}^0$, we let $g'(x)$ denote the lattice path obtained from x by flipping the $d_0(x)$ -th \searrow -step of x touching the line $y = 0$ (in the definition of $g(x)$, we instead flip the $(d_0(x) + 1)$ -th \searrow -step). Similarly, for any $x \in L_{2k,k+1}$, we let $h'(x)$ denote the lattice paths obtained from x by flipping the $(u_1(x) + 1)$ -th \nearrow -step of x touching the line $y = 1$ (in the definition of $h(x)$, we instead flip the $u_1(x)$ -th \nearrow -step). For any of the four mappings $\gamma \in \{g, g', h, h'\}$ from before and any lattice path x from the corresponding domain, we let $\text{pos}(\gamma, x) \in \{1, 2, \dots, 2k\}$ denote the position in which the lattice paths x and $\gamma(x)$ differ.

Note that g is well-defined, as a lattice path $x \in L_{2k,k} \setminus D_{2k}^k$ has less than k flaws and therefore at least one \searrow -step that ends on the line $y = 0$ (which does not contribute to $d_0(x)$), so x has at least $(d_0(x) + 1)$ many \searrow -steps touching the line $y = 0$. Similarly, h' is well-defined, as a lattice path $x \in L_{2k,k+1}$ starts at $(0, 0)$ and ends at $(2k, 2)$ and therefore has at least one \nearrow -step that ends on the line $y = 1$ (which does not contribute to $u_1(x)$), so x has at least $(u_1(x) + 1)$ many \nearrow -steps touching the line $y = 1$. One can argue along very similar lines to show that g' and h are also well-defined.

Recall from Section 1.1 that the function f is defined for any lattice path $x \in L_{2k,k} \setminus D_{2k}^k$ as $f(x) := h(g(x))$.

Our first two lemmas assert that all the four mappings defined before are bijections, and identify the two pairs of mappings that are inverse to each other.

Lemma 6. *The mappings g and h' are bijections between the sets $L_{2k,k} \setminus D_{2k}^k$ and $L_{2k,k+1}$ and we have $g^{-1} = h'$.*

Proof. Consider the lattice path $x \in L_{2k,k} \setminus D_{2k}^k$, its image $x' := g(x)$ and the position $i := \text{pos}(g, x)$ where they differ. We let x_ℓ and x_r , respectively, denote the subpaths of x strictly to the left and to the right of position i . Similarly, we let x'_ℓ and x'_r , respectively, denote the subpaths of x' strictly to the left and to the right of position i . We clearly have

$$x'_\ell = x_\ell \quad , \quad (1)$$

and as the lattice path x'_r is obtained by shifting x_r up by two units we have

$$u_1(x'_r) = u_{-1}(x_r) \quad . \quad (2)$$

By the definition of g we have

$$d_0(x) = d_0(x_\ell) + d_1(x_\ell) \quad . \quad (3)$$

We distinguish the cases whether the \searrow -step at position i of x starts on the line $y = 0$ or on the line $y = 1$. We first consider the case that this \searrow -step starts on the line $y = 0$. In this case we have

$$d_0(x) = d_0(x_\ell) + d_0(x_r) + 1 \quad , \quad (4)$$

$$u_1(x') = u_1(x'_\ell) + u_1(x'_r) \quad . \quad (5)$$

It is also easy to see that

$$u_0(x_\ell) = d_1(x_\ell) \quad , \quad (6)$$

$$u_{-1}(x_r) = d_0(x_r) + 1 \quad . \quad (7)$$

Combining our previous observations we obtain

$$u_1(x'_r) \stackrel{(2),(7)}{=} d_0(x_r) + 1 \stackrel{(3),(4)}{=} d_1(x_\ell) \stackrel{(6),(1)}{=} u_0(x'_\ell) \quad (8)$$

and therefore

$$u_1(x') \stackrel{(5),(8)}{=} u_1(x'_\ell) + u_0(x'_\ell) \quad . \quad (9)$$

The reasoning in the case where the \searrow -step at position i of x starts on the line $y = 1$ can be obtained by straightforward modifications of the above equations. Specifically, we have to subtract 1 on the right-hand sides of (4) and (7), and add 1 on the right-hand sides of (5) and (6), yielding the same resulting relation (9). We obtain from (9) that the position of the $(u_1(x') + 1)$ -th \nearrow -step of x' touching the line $y = 1$ equals the position $\text{pos}(g, x) = i$ of the \searrow -step in x that is flipped by g . Using the definition of h' , it follows that $g^{-1} = h'$. As both sets $L_{2k,k} \setminus D_{2k}^k$ and $L_{2k,k+1}$ have the same size $\binom{2k}{k} - \frac{1}{k+1}\binom{2k}{k} = \binom{2k}{k+1}$, the mappings g and h' are indeed bijections. \square

Lemma 7. *The mappings h and g' are bijections between the sets $L_{2k,k+1}$ and $L_{2k,k} \setminus D_{2k}^0$ and we have $h^{-1} = g'$.*

Proof. This proof is obtained by small modifications to the proof of Lemma 6. As $x' := g'(x)$ is obtained by flipping the $d_0(x)$ -th \searrow -step of x , we have to add 1 on the right-hand side of (3). Equations (1), (2), (4)–(7) remain the same in their respective cases so that the resulting equation (9) changes to

$$u_1(x') = u_1(x'_\ell) + u_0(x'_\ell) + 1 .$$

We conclude from this relation that the position of the $u_1(x')$ -th \nearrow -step of x' touching the line $y = 1$ equals the position $\text{pos}(g', x)$ of the \searrow -step in x that is flipped by g' . Using the definition of h , it follows that $h^{-1} = g'$. As both sets $L_{2k,k+1}$ and $L_{2k,k} \setminus D_{2k}^0$ have the same size, we obtain that h and g' are indeed bijections. \square

The next two lemmas show that the two positions that are flipped when applying g and h consecutively (in one of the two orders) are as close to each other as possible.

Lemma 8. *For any lattice path $x \in L_{2k,k} \setminus D_{2k}^k$ we have $\text{pos}(h, g(x)) < \text{pos}(g, x)$ and in x and $g(x)$ there is no \nearrow -step touching the line $y = 1$ at any position i with $\text{pos}(h, g(x)) < i < \text{pos}(g, x)$.*

Proof. Let $x \in L_{2k,k} \setminus D_{2k}^k$ and $x' := g(x)$. By Lemma 6 we know that h' is the inverse mapping of g . This means that the $(u_1(x') + 1)$ -th \nearrow -step of x' touching the line $y = 1$ is at position $\text{pos}(g, x)$. The mapping h however already flips the $u_1(x')$ -th \nearrow -step of x' touching the line $y = 1$ (the previous one with that property), and the lemma follows. \square

Lemma 9. *For any lattice path $x \in L_{2k,k+1}$ we have $\text{pos}(h, x) < \text{pos}(g, h(x))$ and in $h(x)$ there is no \searrow -step touching the line $y = 0$ at any position i with $\text{pos}(h, x) < i < \text{pos}(g, h(x))$.*

Proof. Let $x \in L_{2k,k+1}$ and $x' := h(x)$. By Lemma 7 we know that g' is the inverse mapping of h . This means that the $d_0(x')$ -th \searrow -step of x' touching the line $y = 0$ is at position $\text{pos}(h, x)$. The mapping g however flips the $(d_0(x') + 1)$ -th \searrow -step of x' touching the line $y = 0$ (the next one with that property), and the lemma follows. \square

For any Dyck path $x \in D_{2k}^e$, $e < k$, let $b := \text{pos}(g, x)$ and $a := \text{pos}(h, g(x))$, respectively, be the positions of the \searrow -step and \nearrow -step that are flipped when applying f . We refer to the subpath of x strictly to the left of position a as the *initial part*, to the subpath strictly between positions a and b as the *middle part*, and to the subpath strictly to the right of position b as the *terminal part*. As mentioned before, $f(x)$ is obtained from x by flipping the steps at positions a and b , leaving the initial and terminal part of x unchanged, and shifting down the middle part down by two units. In particular, no step of x is shifted up by f .

The next lemma shows that our minimum-change bijection increases the number of flaws by exactly 1.

Lemma 10. *For any $x \in D_{2k}^e$, $e < k$, we have that $f(x) \in D_{2k}^{e+1}$.*

Proof. Each Dyck path with e flaws has exactly e many \searrow -steps and exactly e many \nearrow -steps below the line $y = 0$. To prove the lemma it therefore suffices to show that for each $x \in D_{2k}^e$, the image $f(x)$ has exactly $e + 1$ many \nearrow -steps below the line $y = 0$.

First note that any \nearrow -steps of x below the line $y = 0$ is also below the line $y = 0$ in $f(x)$: This is because the \nearrow -step of x flipped by f (it touches the line $y = 1$ by definition) lies above the line $y = 0$ (in x), and because f does not shift up any steps. In contrast to that, the \searrow -step of x that is flipped by f (it touches the line $y = 0$ by definition) creates an additional \nearrow -step of $f(x)$ below the line $y = 0$. It remains to argue that shifting down the middle part of x does not create additional \nearrow -steps below the line $y = 0$. To see this note that such an \nearrow -step in the middle part of x would touch the line $y = 1$ (in x before the shift), which is impossible because of Lemma 8. This completes the proof. \square

The following lemma is illustrated in Figure 9 below. It allows us to recognize for a given Dyck path $x \in D_{2k}^e$ the corresponding preimage $x^0 \in D_{2k}^0$ such that $f^e(x^0) = x$, and will be used later in our algorithm that computes repeated applications of f efficiently.

Lemma 11. *For any $x \in D_{2k}^e$, let U_x be the set of all positions of \nearrow -steps of x that lie below the line $y = 0$, and let D_x be the set of all positions of \searrow -steps that lie below the line $y = -1$ or that are among the first $d_0(x)$ many \searrow -steps touching the line $y = 0$. Moreover, let $x^0 \in D_{2k}^0$ be such that $f^e(x^0) = x$. Then the Dyck paths x^0 and x differ exactly in the positions $U_x \cup D_x$ (in particular, $|U_x| = |D_x| = e$).*

Proof. We argue by induction over the number of flaws. The lemma clearly holds for any $x \in D_{2k}^0$ (in this case $U_x = D_x = \emptyset$), settling the induction basis. For the induction step we fix some $x \in D_{2k}^e$, $e < k$, for which the statement holds, and show that it also holds for $x' := f(x)$.

Let $x^0 \in D_{2k}^0$ be such that $f^e(x^0) = x$. Moreover, let a and b , respectively, be the positions of the \nearrow -step and \searrow -step of $x = (x_1, x_2, \dots, x_{2k})$ that are flipped by f . By the definition of f we know that x_b is the $(d_0(x) + 1)$ -th \searrow -step of x touching the line $y = 0$, and hence $b \notin D_x$, i.e., both x and x^0 have a \searrow -step at position b . As argued in the proof of Lemma 10, when applying f to x the downstep x_b of x is the only additional \nearrow -step of x' below the line $y = 0$ (all others remain the same), implying that $U_{x'} = U_x \cup \{b\}$. By induction we know that U_x contains the positions of all \nearrow -steps of x where x^0 has a \searrow -step. Using that both x and x^0 have a \searrow -step at position b , it follows that $U_{x'}$ contains the positions of all \nearrow -steps of x' where x^0 has a \searrow -step.

We proceed to show that $D_{x'} = D_x \cup \{a\}$ by showing both directions of this inclusion. We begin by showing that $D_x \cup \{a\} \subseteq D_{x'}$. Fix an element $c \in D_x$. If the \searrow -step x_c at position c of x lies below the line $y = -1$, then it is not flipped by f , and also not shifted up, implying that $c \in D_{x'}$ in this case. Otherwise, the \searrow -step x_c is among the first $d_0(x)$ many \searrow -steps of x touching the line $y = 0$. If x_c belongs to the middle part of x , then it is shifted down by two units when applying f , so the corresponding \searrow -step x'_c lies below the line $y = -1$, yielding $c \in D_{x'}$. It remains to consider the case that x_c belongs to the initial part of x . By Lemma 7 and the definition of g' , the \searrow -step x'_a (which is an \nearrow -step in x that is flipped by f) is the $d_0(x')$ -th \searrow -step of x' touching the line $y = 0$ (*), implying that the \searrow -step x'_c is among the first $d_0(x')$ many \searrow -steps of x' , i.e., we have $c \in D_{x'}$. From (*) we also conclude that $a \in D_{x'}$. We therefore obtain $D_x \cup \{a\} \subseteq D_{x'}$, as desired.

We now show that $D_{x'} \subseteq D_x \cup \{a\}$. Let $c \in D_{x'} \setminus \{a\}$. If the corresponding \searrow -step x'_c lies below the line $y = -1$, then the corresponding \searrow -step x_c either also lies below the line $y = -1$, in which case we clearly have $c \in D_x$, or the \searrow -step x_c belongs to the middle part of x that is shifted down when applying f and touches the line $y = 0$ in x . By the definition of f we have $c \in D_x$ also in this case. We now consider the case that the \searrow -step x'_c lies above the line $y = -1$ in x' . Then x'_c belongs to the first $d_0(x')$ many \searrow -steps of x' touching the line $y = 0$. Using Lemma 7 and the definition of g' ,

it follows that the corresponding \searrow -step x_c belongs to the initial part of x and is therefore among the first $d_0(x)$ many \searrow -steps of x touching the line $y = 0$. We therefore have $c \in D_x$ in this case as well. Altogether, we conclude that $D_{x'} = D_x \cup \{a\}$, as claimed.

It remains to verify that $D_{x'}$ contains the positions of all \searrow -steps of x' where x^0 has an \nearrow -step. By induction we know that D_x contains the positions of all \searrow -steps of x where x^0 has an \nearrow -step. As applying f introduces only one additional \searrow -step at position a , we only need to check that x^0 has an \nearrow -step at position a . Recall that the \nearrow -step x_a of x touches the line $y = 1$. It follows that $a \notin U_x$, and so we obtain by induction that x^0 indeed has an \nearrow -step at position a . This completes the proof. \square

We now consider the sequence of flipped positions when applying the minimum-change bijection f repeatedly. Specifically, for any $x \in D_{2k}^0$, we define $x^0 := x$ and $y^i := g(x^i)$ and $x^{i+1} := h(y^i)$ for $i = 0, 1, \dots, k-1$ (by the definition of f we have $x^i = f^i(x)$ for $i = 0, 1, \dots, k$) and define the sequence $p(x) := (x^0, y^0, x^1, y^1, \dots, x^{k-1}, y^{k-1}, x^k)$ and the corresponding sequence of flipped positions $\pi(x) := (\text{pos}(g, x^0), \text{pos}(h, y^0), \text{pos}(g, x^1), \text{pos}(h, y^1), \dots, \text{pos}(g, x^{k-1}), \text{pos}(h, y^{k-1}))$. We call a sequence of integers $(\alpha_1, \alpha_2, \dots, \alpha_{2k})$ *alternating* if we have $\alpha_{i-1} > \alpha_i$ for all $i \in \{2, 4, \dots, 2k\}$ and $\alpha_{i-1} < \alpha_i$ for all $i \in \{3, 5, \dots, 2k-1\}$. For any lattice path $x \in L_{2k,k}$, we denote by \bar{x} the lattice path obtained by mirroring x at the line $y = 0$ (note that $x \in D_{2k}^e$ if and only if $\bar{x} \in D_{2k}^{k-e}$).

The following lemma shows that when applying f repeatedly to some Dyck path with zero flaws, then the order of flipped positions alternate and each step of the path will be flipped exactly once.

Lemma 12. *For any $x \in D_{2k}^0$, the sequence $\pi(x)$ is an alternating permutation of $\{1, 2, \dots, 2k\}$. In particular, we have $f^k(x) = \bar{x}$.*

The alternating permutations $\pi(x)$, $x \in D_{2k}^0$, are shown at the bottom of Figure 4 for $k = 3$. Note also that the Dyck paths in the bottom row of the figure (they have k flaws) are mirror images of the Dyck paths in the top row (they have zero flaws).

Proof. Applying Lemma 11 to $x' := f^k(x) \in D_{2k}^k$ shows that x and x' differ in all $2k$ positions (the sets $U_{x'}$ and $D_{x'}$ satisfy $U_{x'} \cup D_{x'} = \{1, 2, \dots, 2k\}$), implying that $f^k(x) = \bar{x}$. As applying f^k to x flips exactly $2k$ steps in total, we conclude that $\pi(x)$ is indeed a permutation of $\{1, 2, \dots, 2k\}$. Moreover, using Lemma 8 and Lemma 9 we obtain that $\pi(x)$ is alternating. \square

The next lemma allows us to compute the entire sequence $\pi(x)$ recursively directly from x , and will be used later in our algorithm that computes repeated applications of the minimum-change bijection f efficiently. To state the lemma, we define for any sequence $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ the reverse sequence $\text{rev}(\alpha) := (\alpha_n, \alpha_{n-1}, \dots, \alpha_1)$ and we denote the length of the sequence by $|\alpha| := n$. Moreover, for any integer c we define $\alpha + c = c + \alpha := (\alpha_1 + c, \alpha_2 + c, \dots, \alpha_n + c)$. Given a Dyck path $x \in D_{2k}^e$, we may reverse and complement the sequence of steps of x , yielding the Dyck path $\bar{\text{rev}}(x) \in D_{2k}^e$, which is obtained by mirroring x at the vertical line with abscissa k . Given a Dyck path $x \in D_{2k}^0$, let b be the position of the first \searrow -step of $x = (x_1, x_2, \dots, x_{2k})$ touching the line $y = 0$, and define $u := (x_2, \dots, x_{b-1})$ and $v := (x_{b+1}, \dots, x_{2k})$. We refer to this decomposition $x = \nearrow \circ u \circ \searrow \circ v$ as the *canonical decomposition* of x .

Lemma 13. *Given $x \in D_{2k}^0$, let $x = \nearrow \circ u \circ \searrow \circ v$ be the canonical decomposition of x . Then the sequence $\pi(x)$ satisfies the relation $\pi(x) = (|u| + 2, |u| + 2 - \pi(\bar{\text{rev}}(u)), 1, |u| + 2 + \pi(v))$.*

Lemma 13 is illustrated in Figure 7. Intuitively, for any Dyck path $x \in D_{2k}^0$ the sequence $\pi(x)$ can be computed by considering all the ‘hills’ of x that are located at an even height from left to right and those located at an odd height from right to left (in each recursion step, the direction changes),




FIGURE 7. Illustration of Lemma 13. The gray line shows the order in which steps of the Dyck path $x \in D_{2k}^0$ are flipped.

and by flipping the steps in each ‘hill’ in reverse order (first the last step, then the part in the middle recursively, then the first step).

Lemma 13 is an immediate consequence of the following more general, but slightly more technical lemma. To state the lemma we introduce some more definitions.

Given a Dyck path $\hat{x} \in D_{2k}^e$, we refer to a subpath x of \hat{x} with 2ℓ steps that starts and ends on some line $y = c$ and that has no \searrow -steps below the line $y = c$ as a *Dyck subpath of \hat{x}* . We use this wording to refer to a subpath of a surrounding larger Dyck path where the subpath is *not* shifted to the origin. Of course, shifting the subpath x to the origin, x can also be viewed as Dyck path from $D_{2\ell}^0$. This interpretation is adopted whenever functions such as $\overline{\text{rev}}$ or π are applied to x . Moreover, let x be a Dyck subpath of \hat{x} at positions $a+1, a+2, \dots, a+2\ell$, and let σ be the subsequence of $\pi(\hat{x})$ consisting of exactly the elements from $\{a+1, a+2, \dots, a+2\ell\}$. Then we define $\pi(\hat{x})|_x := \sigma - a$ (by Lemma 12 the sequence $\pi(\hat{x})|_x$ is a permutation on the set $\{1, 2, \dots, 2\ell\}$).

Lemma 14. *Let $\hat{x} \in D_{2k}^0$ and $c \in \mathbb{N}$, and let x be a Dyck subpath of \hat{x} at positions $a+1, a+2, \dots, a+2\ell$ that starts and ends on the line $y = c$. Then the elements from $\{a+1, a+2, \dots, a+2\ell\}$ appear consecutively in $\pi(\hat{x})$. Moreover, if c is even we have $\pi(\hat{x})|_x = \pi(x) = (|u|+2, |u|+2 - \pi(\overline{\text{rev}}(u)), 1, |u|+2 + \pi(v))$, where $x = \nearrow \circ u \circ \searrow \circ v$ is the canonical decomposition of $x \in D_{2\ell}^0$. On the other hand, if c is odd we have $\pi(\hat{x})|_x = |x| + 1 - \pi(\overline{\text{rev}}(x))$.*

Proof. For the reader’s convenience, the notations used in this proof are illustrated in Figure 8.

Let $\hat{x} = (\hat{x}_1, \hat{x}_2, \dots, \hat{x}_{2k}) \in D_{2k}^0$ and $x = (\hat{x}_{a+1}, \hat{x}_{a+2}, \dots, \hat{x}_{a+2\ell})$ be as in the lemma. We prove the lemma by induction on $\ell = |x|/2$. The statement clearly holds if $\ell = 0$, settling the induction basis. For the induction step we assume that $\ell \geq 1$ and that the statement holds for all Dyck subpaths of \hat{x} of length strictly less than 2ℓ , and prove that it also holds for x .

We first consider the case that c is even. For the canonical decomposition $x = \nearrow \circ u \circ \searrow \circ v$ of $x \in D_{2\ell}^0$, we consider the Dyck subpaths u and v of x in \hat{x} (see Figure 8 (a)).

In the following we examine the sequence of Dyck paths $\hat{x}, f(\hat{x}), \dots, f^k(\hat{x})$ and in which order the steps of the subpath x are flipped when repeatedly applying f (recall from Lemma 12 that each step of \hat{x} and x is flipped exactly once). Let r be the largest integer such that the positions flipped during the first $2r$ applications of f are not in the set $\{a+1, a+2, \dots, a+2\ell\}$ (equivalently, the first $2r$ entries of $\pi(\hat{x})$ are not from this set). By definition, in $f^r(\hat{x})$ none of the steps at positions $a+1, a+2, \dots, a+2\ell$ are flipped yet (w.r.t. \hat{x}), so the subpath x' of $f^r(\hat{x})$ at these positions is simply a shifted-down copy of x (in particular, x' is a Dyck subpath of $f^r(\hat{x})$). Let u' and v' , respectively, be the corresponding shifted-down copies of u and v (see Figure 8 (b)).

We claim that the Dyck subpath x' of $f^r(\hat{x})$ starts and ends on the line $y = 0$ (see Figure 8 (b)). To see this note that x must be completely contained either in the initial, middle, or terminal part

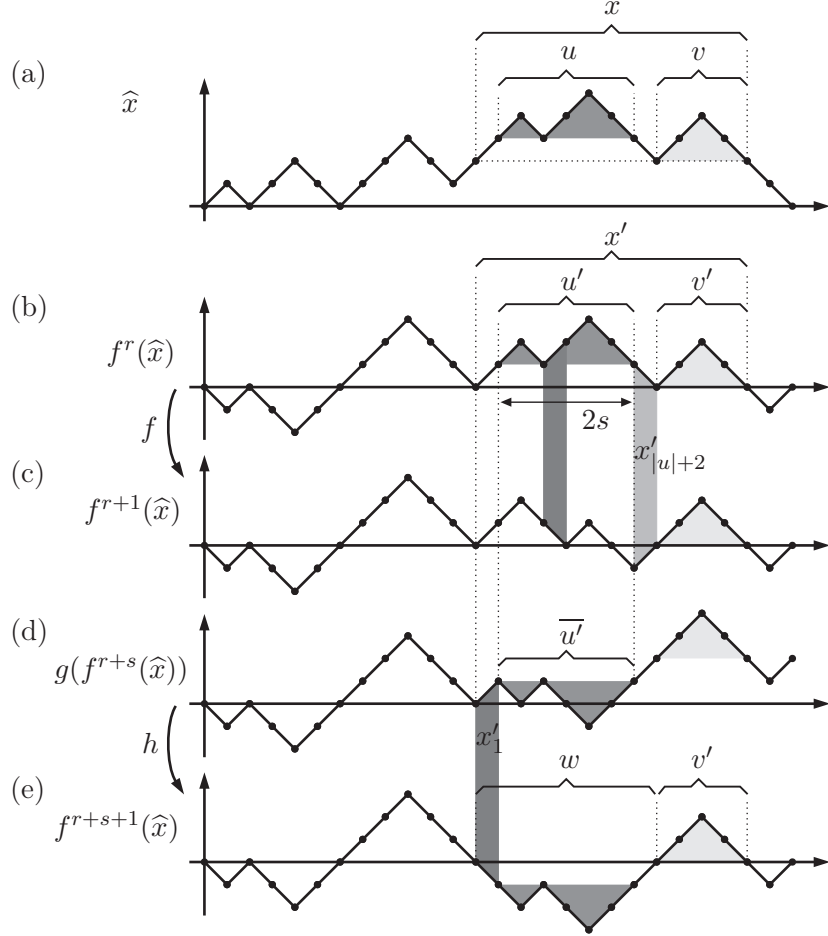


FIGURE 8. Illustration of the notations used in the proof of Lemma 14.

of each of the Dyck paths $\hat{x}, f(\hat{x}), \dots, f^{r-1}(\hat{x})$, and that this subpath is shifted down by two units if it is contained the middle part and not shifted otherwise. As c is even, we obtain that in each of the Dyck paths $\hat{x}, f(\hat{x}), \dots, f^r(\hat{x})$ the subpath x starts and ends on a horizontal line at an even height. Let c' be this even height with respect to $f^r(\hat{x})$. Note that c' cannot be negative as otherwise the first step of x' (an \nearrow -step) is an unflipped \nearrow -step below the line $y = 0$ in $f^r(\hat{x})$, contradicting Lemma 11. Moreover, as f does not flip any steps above the line $y = 2$ we conclude that $c' = 0$, as claimed.

Next we show that applying f to $f^r(\hat{x})$ flips a \searrow -step of x' . For the sake of contradiction assume that an \nearrow -step of x' and a \searrow -step strictly to the right of position $a + 2\ell$ are flipped. Then in $f^{r+1}(\hat{x})$ the \searrow -step at position $a + 2\ell$ (obtained by shifting down the last step of x' by two units) is unflipped (w.r.t. \hat{x}) and lies below the line $y = -1$, contradicting Lemma 11. Therefore, applying f to $f^r(\hat{x})$ indeed flips a \searrow -step of x' .

We proceed to show that when applying f to $f^r(\hat{x})$ the \searrow -step of $x' = (x'_1, x'_2, \dots, x'_{2\ell})$ being flipped is exactly the step $x'_{|u|+2}$ (see Figure 8 (b) and (c)). Clearly, the \searrow -step being flipped touches the line $y = 0$ and the step $x'_{|u|+2}$ is the leftmost one in x' with that property. Suppose for the sake of contradiction that the \searrow -step $x'_{|u|+2}$ is not flipped, but instead a \searrow -step of x' further to the right, then by the definition of f the step $x'_{|u|+2}$ would be among the first $d_0(f^r(\hat{x}))$ many \searrow -steps of $f^r(\hat{x})$

touching the line $y = 0$. This however contradicts Lemma 11, from which it would follow that $x'_{|u|+2}$ is flipped in $f^r(\hat{x})$ (w.r.t. \hat{x}).

Consequently, applying f to $f^r(\hat{x})$ flips the \searrow -step $x'_{|u|+2}$ and by Lemma 8 the first \nearrow -step of $f^r(\hat{x})$ touching the line $y = 1$ to the left of it. As the \nearrow -step x'_1 touches the line $y = 1$, the \nearrow -step being flipped is part of u' if u' is non-empty, and equal to x'_1 otherwise.

Note that u is a Dyck subpath of x that starts and ends on the line $y = c + 1$ and has length strictly less than 2ℓ . We therefore obtain by induction that $\pi(\hat{x})|_u = |u| + 1 - \pi(\overline{\text{rev}}(u))$. In particular, while no steps of u' are flipped yet in $g(f^r(\hat{x}))$ (w.r.t. \hat{x}), all steps of u' are flipped in $g(f^{r+s}(x))$, where $s := |u|/2$ (see Figure 8 (d)).

Now it is easy to see that x'_1 is flipped when applying h to $g(f^{r+s}(\hat{x}))$ (see Figure 8 (d) and (e)). Indeed, in $g(f^{r+s}(x))$ the \nearrow -step x'_1 is the first one touching the line $y = 1$ to the left of $\overline{u'}$, so by Lemma 8 this step is flipped by h .

Let us now consider the Dyck path $h(g(f^{r+s}(x))) = f^{r+s+1}(x)$ (see Figure 8 (e)), and let w be the subpath at positions $a + 1, a + 2, \dots, a + |u| + 2$ (we have $w = \searrow \circ \overline{u'} \circ \nearrow$). By Lemma 9, applying f to $f^{r+s+1}(x)$ flips the first \searrow -step touching the line $y = 0$ to the right of w , which is clearly a \searrow -step of v' if v' is non-empty. As the length of v' is strictly less than 2ℓ , we obtain by induction that $\pi(\hat{x})|_v = \pi(v)$. In particular, when repeatedly applying f to $f^{r+s+1}(x)$, first all steps of v' are flipped, and then all other unflipped (w.r.t. \hat{x}) steps of $f^{r+s+1}(x)$.

Combining our previous observations we obtain $\pi(\hat{x})|_x = (|u| + 2, |u| + 2 - \pi(\overline{\text{rev}}(u)), 1, |u| + 2 + \pi(v))$. Repeating the above argument with the Dyck path $\hat{x} := x$ shows that this expression also equals $\pi(x)$, as claimed.

We now consider the case that c is odd. The analysis in this case is very similar to the case where c is even, so we omit most of the details. For the canonical decomposition $\overline{\text{rev}}(x) = \nearrow \circ u \circ \searrow \circ v$ of $\overline{\text{rev}}(x) \in D_{2\ell}^0$, i.e., we have $x = \overline{\text{rev}}(v) \circ \nearrow \circ \overline{\text{rev}}(u) \circ \searrow$, we consider the Dyck subpaths $\overline{\text{rev}}(v)$ and $\overline{\text{rev}}(u)$ of x in \hat{x} . We let r be the maximal integer such that the positions flipped during the first $2r$ applications of f to \hat{x} are not in the set $\{a + 1, a + 2, \dots, a + 2\ell\}$, and let x' be the subpath of $f^r(\hat{x})$ at these positions (x' is a shifted-down copy of x). Arguing along similar lines as in the previous case, one can show that $x' = (x'_1, x'_2, \dots, x'_{2\ell})$ starts and ends on the line $y = 1$ in $f^r(x)$, and that the \nearrow -step $x'_{|v|+1}$ is flipped when applying f to $f^r(\hat{x})$. Moreover, after flipping $x'_{|v|+1}$, all the steps of the shifted-down copy of $\overline{\text{rev}}(u)$ are flipped in the order $\pi(\overline{\text{rev}}(u))$ (by induction), followed by the \searrow -step $x'_{|u|+|v|+2} = x'_{2\ell}$, followed by all the steps of the shifted-down copy of $\overline{\text{rev}}(v)$ in the order $|v| + 1 - \pi(v)$ (by induction). Combining these observations we obtain

$$\begin{aligned} \pi(\hat{x})|_x &= (|v| + 1, |v| + 1 + \pi(\overline{\text{rev}}(u)), |u| + |v| + 2, |v| + 1 - \pi(v)) \\ &= |x| + 1 - (|u| + 2, |u| + 2 - \pi(\overline{\text{rev}}(u)), 1, |u| + 2 + \pi(v)) \\ &= |x| + 1 - \pi(\overline{\text{rev}}(x)) , \end{aligned}$$

where in the second to last step we used $|x| = |u| + |v| + 2$, and in the last step we used the first part of the lemma. This completes the proof. \square

2.2. Proofs of Theorems 2–5. With the auxiliary lemmas from the previous section in hand, we are now ready to present the proofs of our main results.

Proof of Theorem 2. Combining Lemmas 6, 7 and 10 we obtain that for any $e \in \{0, 1, \dots, k - 1\}$ the mapping f is an injection between D_{2k}^e and D_{2k}^{e+1} . This implies $|D_{2k}^0| \leq |D_{2k}^1| \leq \dots \leq |D_{2k}^k|$, and as D_{2k}^0 and D_{2k}^k have the same size by symmetry, all these inequalities are in fact equalities. In particular, f is a bijection between D_{2k}^e and D_{2k}^{e+1} . \square




FIGURE 9. Illustration of Lemma 11 and of the data structures used in our algorithm from Theorem 3. Note that x^0 and x differ exactly at the positions in $U_x \cup D_x$. The horizontal bidirectional arrows in the upper part of the figure represent the pointer array used for fast navigation along the Dyck path x^0 when computing $\pi(x^0)$ recursively.

Proof of Theorem 3. We begin by describing an algorithm that computes for a given Dyck path $x \in D_{2k}^e$ the sequence of Dyck paths $f(x), f^2(x), \dots, f^{k-e}(x)$, and then argue about the running time and space requirements of the algorithm.

In the initialization phase, for a given Dyck path $x \in D_{2k}^e$, our algorithm first computes the corresponding Dyck path $x^0 \in D_{2k}^0$ such that $f^e(x^0) = x$ with the help of Lemma 11 (see Figure 9). We then compute the sequence $\pi(x^0)$ using the recursive rule described in Lemma 13. With the help of $\pi(x^0)$ we can then compute each of the Dyck paths $f^i(x)$, $i = 1, 2, \dots, k - e$, by flipping in the i -iteration the two steps in x at the positions that are given by the entries of $\pi(x^0)$ at positions $2(e + i) - 1$ and $2(e + i)$ (a \searrow -step and an \nearrow -step, respectively).

Computing x^0 can clearly be achieved in time $\mathcal{O}(k)$, and with $\pi(x^0)$ at hand, each of the Dyck paths $f(x), f^2(x), \dots, f^{k-e}(x)$ can easily be computed in time $\mathcal{O}(1)$. It remains to show how to compute $\pi(x^0)$ recursively in time $\mathcal{O}(k)$ (time spent in the initialization phase). For this we precompute an array of bidirectional pointers below the ‘hills’ of x^0 between neighboring pairs of an \nearrow -step and a \searrow -step on the same height level (see the upper part of Figure 9). With the help of these pointers, the canonical decomposition of x^0 into subpaths u and v required in Lemma 13 can be performed in constant time. Moreover, throughout the recursive computation of $\pi(x^0)$ we maintain the two boundary indices of the subpath currently under consideration and whether this subpath is considered in forward or backward direction (the direction alternates in each recursion step), so that the flip positions in the recursion formula can be computed directly from these boundary indices in constant time (recall the remarks after Lemma 13 and see the comments of our C++ implementation for details [www]). It follows that the overall running time for the recursion is $\mathcal{O}(k)$, as desired.

The space required by all the above-mentioned data structures is clearly $\mathcal{O}(k)$. This completes the proof. \square

Proof of Theorem 4. For this proof we interpret the lattice paths in $L_{2k,k}$ and $L_{2k,k+1}$ as subsets of the ground set $[2k] := \{1, 2, \dots, 2k\}$, where an \nearrow -step at position i means that the element i is contained in the subset, and a \searrow -step at position i means that the element i is not contained in the subset. For any $x \in D_{2k}^0$ we consider the sequence $p(x) = (x^0, y^0, x^1, y^1, \dots, x^{k-1}, y^{k-1}, x^k)$ defined before Lemma 12. Viewing all the x^i as k -element subsets of $[2k]$, and the y^i as $(k + 1)$ -element subsets of $[2k]$, by the definition of the sequence $p(x)$ we have $y^i \supset x^i$ and $y^i \supset x^{i+1}$ for all




FIGURE 10. Illustration of the notations used in the proof of Theorem 4.

$i = 0, 1, \dots, k-1$. We now define $\bar{y}^i := ([2k] \setminus y^i) \cup \{2k+1\}$ and consider the modified sequence $c(x) := (x^0, \bar{y}^0, x^1, \bar{y}^1, \dots, x^{k-1}, \bar{y}^{k-1}, x^k)$. The sets x^i and \bar{y}^i are k -element subsets of $[2k+1]$, and we have $\bar{y}^i \cap x^i = \emptyset$ and $\bar{y}^i \cap x^{i+1} = \emptyset$ for all $i = 0, 1, \dots, k-1$. Moreover, by Lemma 12 we have $x^k = [2k] \setminus x^0$, so $x^0 \cap x^k = \emptyset$. We can thus interpret $c(x)$ as a cycle of length $2k+1$ in the odd graph O_{2k+1} (see Figure 10). By Lemma 6 and Lemma 7, the cycles $c(x)$, $x \in D_{2k}^0$, are all disjoint, and they must consequently form a \mathcal{C}_{2k+1} -factor of this graph that consists of $|D_{2k}^0| = C_k$ many cycles (the total number of vertices visited by these cycles is $(2k+1)C_k$, which is equal to $\binom{2k+1}{k}$, the total number of vertices of O_{2k+1}). \square

The following simple lemma allows us to transfer cycle-factors with cycles of odd length from the odd graph to the middle levels graph, and will be used to prove Theorem 5.

Lemma 15. *For any $k \geq 1$ and any odd $\ell \geq 3$, if the odd graph O_{2k+1} has a \mathcal{C}_ℓ -factor, then the middle levels graph M_{2k+1} has a $\mathcal{C}_{2\ell}$ -factor.*

Proof. Let m be such that $\ell = 2m+1$, and for any cycle $c = (x^1, x^2, \dots, x^{2m+1})$ of the cycle-factor of the odd graph O_{2k+1} consider the corresponding cycle

$$c' := (x^1, \bar{x}^2, x^3, \bar{x}^4, \dots, x^{2m-1}, \bar{x}^{2m}, x^{2m+1}, \bar{x}^1, x^2, \bar{x}^3, x^4, \dots, \bar{x}^{2m-1}, x^{2m}, \bar{x}^{2m+1})$$

in the middle levels graph M_{2k+1} , where $\bar{x}^i := [2k+1] \setminus x^i$ for all $i = 1, 2, \dots, \ell$. It is easy to see that each of these cycles has length 2ℓ , and that together they form a cycle-factor of M_{2k+1} (the middle levels graph has twice as many vertices as the odd graph). \square

Proof of Theorem 5. Combine Theorem 4 and Lemma 15. \square

REFERENCES

- [BER76] J. Bitner, G. Ehrlich, and E. Reingold. Efficient generation of the binary reflected Gray code and its applications. *Comm. ACM*, 19(9):517–521, 1976.
- [Big79] N. Biggs. Some odd graph theory. In *Second International Conference on Combinatorial Mathematics (New York, 1978)*, volume 319 of *Ann. New York Acad. Sci.*, pages 71–81. New York Acad. Sci., New York, 1979.
- [BR98] B. Bultena and F. Ruskey. An Eades-McKay algorithm for well-formed parentheses strings. *Inform. Process. Lett.*, 68(5):255–259, 1998.
- [BW84] M. Buck and D. Wiedemann. Gray codes with restricted density. *Discrete Math.*, 48(2-3):163–171, 1984.
- [CF49] K. Chung and W. Feller. On fluctuations in coin-tossing. *Proceedings of the National Academy of Sciences of the United States of America*, 35(10):605–608, 1949.
- [CF02] Y. Chen and Z. Füredi. Hamiltonian Kneser graphs. *Combinatorica*, 22(1):147–149, 2002.
- [CGH13] L. Chua, A. Gyárfás, and C. Hossain. Gallai-colorings of triples and 2-factors of B_3 . *Int. J. Comb.*, pages Art. ID 929565, 6 pp., 2013.
- [Cha89] P. Chase. Combination generation and graylex ordering. *Congr. Numer.*, 69:215–242, 1989. Eighteenth Manitoba Conference on Numerical Mathematics and Computing (Winnipeg, MB, 1988).
- [Che00] Y. Chen. Kneser graphs are Hamiltonian for $n \geq 3k$. *J. Combin. Theory Ser. B*, 80(1):69–79, 2000.
- [Che03] Y. Chen. Triangle-free Hamiltonian Kneser graphs. *J. Combin. Theory Ser. B*, 89(1):1–16, 2003.
- [Che08] Y. Chen. The Chung-Feller theorem revisited. *Discrete Math.*, 308(7):1328–1329, 2008.
- [CL87] B. Chen and K. Lih. Hamiltonian uniform subset graphs. *J. Combin. Theory Ser. B*, 42(3):257–263, 1987.

- [DG12] P. Diaconis and R. Graham. *Magical mathematics*. Princeton University Press, Princeton, NJ, 2012. The mathematical ideas that animate great magic tricks, With a foreword by Martin Gardner.
- [DKS94] D. Duffus, H. Kierstead, and H. Snevily. An explicit 1-factorization in the middle of the Boolean lattice. *J. Combin. Theory Ser. A*, 65(2):334–342, 1994.
- [DSW88] D. Duffus, B. Sands, and R. Woodrow. Lexicographic matchings cannot form Hamiltonian cycles. *Order*, 5(2):149–161, 1988.
- [EFY05] S. Eu, T. Fu, and Y. Yeh. Refined Chung-Feller theorems for lattice paths. *J. Combin. Theory Ser. A*, 112(1):143–162, 2005.
- [Ehr73] G. Ehrlich. Loopless algorithms for generating permutations, combinations, and other combinatorial configurations. *J. Assoc. Comput. Mach.*, 20:500–513, 1973.
- [EHR84] P. Eades, M. Hickey, and R. Read. Some Hamilton paths and a minimal change algorithm. *J. Assoc. Comput. Mach.*, 31(1):19–29, 1984.
- [EM84] P. Eades and B. McKay. An algorithm for generating subsets of fixed size with a strong minimal change property. *Inform. Process. Lett.*, 19(3):131–133, 1984.
- [FT95] S. Felsner and W. Trotter. Colorings of diagrams of interval orders and α -sequences of sets. *Discrete Math.*, 144(1-3):23–31, 1995. Combinatorics of ordered sets (Oberwolfach, 1991).
- [Gra] F. Gray. Pulse code communication. March 17, 1953 (filed Nov. 1947). U.S. Patent 2,632,058.
- [GŠ10] P. Gregor and R. Škrekovski. On generalized middle-level problem. *Inform. Sci.*, 180(12):2448–2457, 2010.
- [Hav83] I. Havel. Semipaths in directed cubes. In *Graphs and other combinatorial topics (Prague, 1982)*, volume 59 of *Teubner-Texte Math.*, pages 101–108. Teubner, Leipzig, 1983.
- [HKRR05] P. Horák, T. Kaiser, M. Rosenfeld, and Z. Ryjáček. The prism over the middle-levels graph is Hamiltonian. *Order*, 22(1):73–81, 2005.
- [Hod55] J. Hodges. Galton’s rank-order test. *Biometrika*, 42(1/2):261–262, 1955.
- [HW78] K. Heinrich and W. Wallis. Hamiltonian cycles in certain graphs. *J. Austral. Math. Soc. Ser. A*, 26(1):89–98, 1978.
- [JK04] R. Johnson and H. Kierstead. Explicit 2-factorisations of the odd graph. *Order*, 21(1):19–27 (2005), 2004.
- [JM95] T. Jenkyns and D. McCarthy. Generating all k -subsets of $\{1 \cdots n\}$ with minimal changes. *Ars Combin.*, 40:153–159, 1995.
- [Joh04] R. Johnson. Long cycles in the middle two layers of the discrete cube. *J. Combin. Theory Ser. A*, 105(2):255–271, 2004.
- [Joh11] R. Johnson. An inductive construction for Hamilton cycles in Kneser graphs. *Electron. J. Combin.*, 18(1):Paper 189, 12, 2011.
- [Knu11] D. Knuth. *The Art of Computer Programming, Volume 4A*. Addison-Wesley, 2011.
- [KT88] H. Kierstead and W. Trotter. Explicit matchings in the middle levels of the Boolean lattice. *Order*, 5(2):163–171, 1988.
- [Mac09] P. MacMahon. Memoir on the theory of the partitions of numbers. Part iv. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 209(441-458):153–175, 1909.
- [Mat76] M. Mather. The Rugby footballers of Croam. *J. Combinatorial Theory Ser. B*, 20(1):62–63, 1976.
- [ML72] G. Meredith and K. Lloyd. The Hamiltonian graphs O_4 to O_7 . In *Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972)*, pages 229–236. Inst. Math. Appl., Southend-on-Sea, 1972.
- [MN15] T. Mütze and J. Nummenpalo. Efficient computation of middle levels Gray codes. In *Algorithms—ESA 2015*, volume 9294 of *Lect. Notes Comp. Sci.*, pages 915–927. Springer, Heidelberg, 2015.
- [MN16] T. Mütze and J. Nummenpalo. A constant-time algorithm for middle levels Gray codes. *arXiv:1606.06172*. An extended abstract has been presented at SODA 2017., Jul 2016.
- [MS15] T. Mütze and P. Su. Bipartite Kneser graphs are Hamiltonian. *Electr. Notes Discr. Math.*, 49:259 – 267, 2015. Proceedings of Eurocomb 2015. The full version is available at *arXiv:1503.09175* and will appear in *Combinatorica*.
- [Müt16] T. Mütze. Proof of the middle levels conjecture. *Proc. London Math. Soc.*, 112(4):677–713, 2016.
- [MW12] T. Mütze and F. Weber. Construction of 2-factors in the middle layer of the discrete cube. *J. Combin. Theory Ser. A*, 119(8):1832–1855, 2012.
- [MY09] J. Ma and Y. Yeh. Generalizations of Chung-Feller theorems. *Bull. Inst. Math. Acad. Sin. (N.S.)*, 4(3):299–332, 2009.
- [Nar67] T. Narayana. Cyclic permutation of lattice paths and the Chung-Feller theorem. *Skand. Aktuarietidskr*, 1967:23–30, 1967.
- [NW75] A. Nijenhuis and H. Wilf. *Combinatorial algorithms*. Academic Press, New York-London, 1975. Computer Science and Applied Mathematics.
- [RP90] F. Ruskey and A. Proskurowski. Generating binary trees by transpositions. *J. Algorithms*, 11(1):68–84, 1990.

- [Ruk11] J. Rukavicka. On generalized Dyck paths. *Electron. J. Combin.*, 18(1):Paper 40, 3 pp., 2011.
- [Rus88] F. Ruskey. Adjacent interchange generation of combinations. *J. Algorithms*, 9(2):162–180, 1988.
- [SA11] M. Shimada and K. Amano. A note on the middle levels conjecture. *arXiv:0912.4564*, Sep 2011.
- [Sav93] C. Savage. Long cycles in the middle two levels of the Boolean lattice. *Ars Combin.*, 35-A:97–108, 1993.
- [Sav97] C. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997.
- [SSS09] I. Shields, B. Shields, and C. Savage. An update on the middle levels problem. *Discrete Math.*, 309(17):5271–5277, 2009.
- [Sta15] R. Stanley. *Catalan numbers*. Cambridge University Press, Cambridge, 2015.
- [SW95] C. Savage and P. Winkler. Monotone Gray codes and the middle levels problem. *J. Combin. Theory Ser. A*, 70(2):230–248, 1995.
- [TL73] D. Tang and C. Liu. Distance-2 cyclic chaining of constant-weight codes. *IEEE Trans. Computers*, C-22:176–180, 1973.
- [Wal98] T. Walsh. Generation of well-formed parenthesis strings in constant worst-case time. *J. Algorithms*, 29(1):165–173, 1998.
- [Wil89] H. Wilf. *Combinatorial algorithms: an update*, volume 55 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1989.
- [Win04] P. Winkler. *Mathematical puzzles: a connoisseur’s collection*. A K Peters, Ltd., Natick, MA, 2004.
- [Woa01] W. Woan. Uniform partitions of lattice paths and Chung-Feller generalizations. *Amer. Math. Monthly*, 108(6):556–559, 2001.
- [www] currently <http://www.math.tu-berlin.de/~muetze> and <http://www.math.tu-berlin.de/~wiechert>.